Re: Docket Number 17-108
Restoring Internet Freedom

Benjamin Kreuter

July 17, 2017

1 Summary

As a computer scientist and engineer working for a Web company, I oppose the FCC's proposal to reclassify ISPs as information services on technical and policy grounds. The importance of network neutrality as a principle is beyond dispute: adherence to net neutrality principles has enabled the Internet to be an engine of innovation. Network neutrality is the guiding principle of the Internet's technical standards.

Ensuring the continued adherence to this fundamental principle requires appropriate regulatory measures. In general those measures can take one of two forms: encouraging competitive markets through line-sharing and pole access rules, or the current approach of explicitly forstalling violations. I do not oppose either approach, but my understanding is that Title II classification is required either way. It is also our understanding that Title II classification has been upheld by the courts. Iurge the FCC to retain the current Title II classification, and to work within that framework on improving regulations in pursuit of the public interest.

Any change to network neutrality rules should be predicated on technical, economic, and legal facts. As an engineer I am best able to comment on the technical analysis given in the NPRM. Unfortunately it appears that the NPRM's technical analysis is deeply flawed and based on a confused understanding of the Internet. I have also identified numerous inaccuracies in the NPRM's description of the history of ISP regulation and the history of network neutrality regulations. I believe that these flaws are fatal to the analysis and that the proposed rule changes are tainted by these factual errors.

I believe that politics and political considerations should play *no* role in this issue. I am therefore disturbed by the NPRM's apparent suggestion that the move to Title II classification was a result of a request by President Obama. My understanding is that the move was prompted by the unprecedented number of comments in support of Title II regulation submitted by the general public in response to a proposed weaker form of regulation under Title I. If indeed those comments played no role, and the move was simply political in nature, I believe it is of paramount importance that the FCC carefully consider all public comments submitted in this current proceeding and that any political considerations be ignored. To that effect I hope it is merely coincidental that this proceeding occurred so soon after the election of President Trump. Unfortunately, recent

comments by Commissioners Pai and O'Reilly that indicate an intention to ignore the majority of comments opposed to this plan, and Chairman Pai's selective public reading of a small handful of comments in support of this plan, leaves me with serious doubts that this proceeding is not politically motivated. I hope to be proved wrong.

2 Introduction

It is hard to overstate the value of the Internet. Some of the largest and most innovative companies in America and in the world are Internet-based companies. Students now depend on Internet access to complete their schoolwork. People use the Internet to shop, for entertainment, to manage their finances, to plan travel, and to work. Perhaps most important of all, the Internet has become our primary method of communication, and it is rapidly subsuming every previous communication system.

The impact of the Internet is not mere coincidence. As a system, the Internet is designed to be as general as possible and to support as many applications as its users can invent. The Internet was originally designed to connect research networks together, and to support the various research projects being conducted on those networks. As a result of its generality and ability to support so many applications and to connect a diversity of networks and computers, the Internet rapidly subsumed other computer networks, resulting in the unified global networkIdepend on today.

Network neutrality is the principle that makes the Internet's generality possible. This is explained in detail in section 4, but at a high-level, network neutrality ensures that new applications can be deployed without required complex coordination or negotiation with ISPs. The strongest form of network neutrality is also the simplest: that ISPs should treat all communication across their networks equally, without prioritization and without "intelligence." Somewhat surprisingly, researchers working on the Internet2 project discovered that this is actually *better* for users [12].

In practice adherence to the strongest form of network neutrality is rare among consumer-focused ISPs. Where competition exists markets have tended to choose weaker forms that allow discrimination among application classes but not between applications or services within a class. In this model consumers choose freely among edge services, without interference from ISPs.

2.1 The Myth of Light-Touch Regulation

The NPRM asserts that it was only in 2015 that the FCC applied Title II regulations to ISPs. This, however, is not an accurate portrayal of even the historical regulation of broadband services, let alone ISPs generally. In fact, until 2005, DSL service was regulated under both Title I and Title II: Title I governed the higher-level Internet access service, while Title II covered the infrastructure over which Internet access service was conveyed. At the time Americans enjoyed a highly competitive market for high-speed service, made possible by the use of Title II authority to impose line-sharing rules.

Beginning in 2005 the FCC reclassified ISP service generally as Title I information services, freeing telephone companies from DSL line-sharing rules. Almost overnight the competitive market Americans had come to enjoy evaporated, and the current era of local broadband monopolies

and duopolies began. As the NPRM notes, within a few years it had become clear that without the pressure of a competitive market, ISPs were abusing their power by violating net neutrality principles, with Comcast's inteference with customer BitTorrent traffic being one of the most egregious, and the FCC attempted to remedy the situation through stricter regulation. Though the FCC initially attempted to use its authority under Title I to impose net neutrality regulations, it was ultimately determined by the courts that Title II authority was required, which motivated the 2015 rule change this current NPRM seeks to undo.

In other words, the era of light-touch regulation the NPRM refers to existed for only a brief period of time and led to unsatisfactory outcomes. The idea that a light-touch approach would allow a market to flourish is simply a myth. The evidence from the history of ISP regulation demonstrates beyond any doubt that the light-touch approach would neither lead to a healthy market nor to high-quality Internet services.

3 The Internet

An immediate problem with the NPRM's analysis is its failure to distinguish the networks comprising the Internet, which includes but is not limited to ISP networks, and the edge services that users interact with via the Internet. This seems to stem from a confused understanding of the Internet's design and purpose.

The Internet is a "network of networks" that are interconnected by special computers called *routers*. The purpose of a router is to facilitate communication across different networks. Computers connected to the Internet, including routers, are identified by numeric "IP addresses."

The Internet is a packet-switched network, meaning that the basic unit of communication is a short message known as a *packet*. When a router receives a packet, it examines the packet's destination IP address and uses that address to decide where the packet should be transmitted. If the packet can be re-transmitted, there are two possibilities. The first is that the router is connected to the same network as the destination, in which case the packet is immediately transmitted to its recipient computer. Otherwise, the router will choose another router to forward the packet to, and that next router will repeat the process [3].

It is this process of router-to-router packet switching, based on a global addressing scheme, that defines the Internet. Everything else, including the various edge services familiar to consumers, is built on that system. A modern broadband ISP sells, as a service, a connection to its routers, which connect its network to the Internet.

3.1 Routing is a Communication Service

The NPRM specifically requests comment on its analysis that the packet routing performed by ISPs does not represent a communication service, because routing decisions are based on the architecture of the network rather than specific user instructions. Ibelieve the NPRM is incorrect in its analysis.

First, note that routing decisions *must* be based on some user-supplied information. A packet being routed must have some destination IP address, and that address is fundamental to routing

decisions. It is true that a large edge service may have many IP addresses, and users are rarely aware of which IP addresses they have exchanged packets with, it is equally true that it is the users' computers that decide which IP addresses to send packets to.

The fact that a user's device assigns destination IP addresses is critical. By analogy, it is the equivalent of a telephone that includes a speed-dial feature that automatically dials a phone number on its user's behalf. A similar analogy applies to edge services sometimes having multiple IP addresses: it is the equivalent of a large organization maintaining multiple phone numbers.

The NPRM notes that the Telecommunications Act of 1996 explicitly excludes from the definition of an information service the use of computers for the management or operation of a network. Routers, utilizing information about the architecture of the network, are simply computers used for the operation of the network itself. I believe it is reasonable to view DNS and caching servers similarly, but as discussed in Section 3.2 below the alternative view would not support the claim that ISPs provide an information service.

3.2 Services Provided by ISPs

The NPRM requests comment on its analysis that ISPs provide their customers with the "capability for generating, acquiring, storing, transforming, processing, retrieving, utilizing, or making available information via telecommunications," which would define ISPs as information services. Its analysis points to the fact that users of ISP services can access social media networks, news websites, search engines, and other edge services.

It is certainly true that all of the "capabilities" information services provide to users are available to ISP customers. To access a social media network or to use a search engine, one must have some connection to the Internet, which is what ISPs provide to their customers. However, it is equally true that ISPs do not actually provide any of the capabilities described in Section 3 of the Telecommunications Act. The NPRM errs in its analysis by suggesting that ISPs are providing such capabilities to their customers, when in fact it is the edge services that provide these capabilities.

I believe the NPRM's suggestion that the fact that *some* ISP is needed to access edge services implies that ISPs are making such capabilities available is erroneous. Equally necessary is electricity service, but it should be clear that electric utilities do not make any such capabilities available to their customers. The FCC should treat ISP service as a *prerequisite* to having capabilities provided *by third parties*.

An important distinction must be made between a service provider that negotiates and contracts with third parties to provide capabilities to its customers and the edge services ISP customers access via the Internet. In the case of contracted service, an ISP *necessarily* coordinates with a third party on its customers' behalf, and the service is not generally accessible without some agreement between the ISP and the third party. Edge services do not require coordination with ISPs and their services are accessible regardless of which particular ISP is used [1]. It is true that in

[1] I am aware of only one prominent exception, which is ESPN's streaming service. ESPN only provides service to users whose ISP has paid a fee to ESPN, and that fee cannot be paid on an individual basis. Ibelieve this exceptional case is irrelevant for the statutory analysis, as it is highly unusual and Internet users can easily evade these restrictions

some cases edge services, especially larger services, coordinate with ISPs of their own accord, but there has never been any need for edge services to do so.

The consequence of this lack of coordination is that ISPs cannot provide their customers with any guarantee that any of the above capabilities will be available. Edge services can be inaccessible for a variety of reasons and may at any time be terminated, without any warning to ISPs or to their customers. The only guarantee an ISP can make is that its customers can use their Internet service to access those edge services they have chosen. In fact, the service provided by an ISP is more general than access to edge services: ISPs provide communication between their customers' computers and other Internet-connected computers, including but not limited to edge services.

3.3 Inter-Protocol Translation is a Communication Service

The NPRM asks for comment on its analysis that ISPs provide an information service by, for example, translation of IPv4 addresses to IPv6 addresses on their customers' behalf. Comment was also requested on whether firewalls should be considered information services. I believe that the analysis is incorrect in both cases.

Translation between IP address ranges or formats is a special function of routers that has been standardized by relevant engineering bodies [2, 11]. Typically, routers forward packets without changing the destination or source IP address. In some cases, however, it is convenient or necessary for a router to rewrite one or both of those addresses before forwarding a packet. For example, consumer routers often have just one connection to the larger Internet, and ISPs will often assign just one IPv4 address. It is possible for many computers to share that single address by setting up a network that uses "private" IP addresses i.e. those that by convention are not routed over the Internet, which the router itself will track and rewrite when it forwards packets.

In the case of IPv6, compatibility with IPv4 motivated the developed of several standards for routers translating between IPv6 and IPv4 addresses reliably. As with the case of sharing a single IPv4 address, this is represents a deviation from the typical behavior of routers and requries addresses to be re-written before packets are forwarded.

In none of these cases of address re-writing is the routing not supporting telecommunication. Although re-writing IP addresses in packets is a deviation from the ideal Internet of simple packet forwarding, practical considerations and limitations in the design of IPv4 have necessitated such approaches.

Packet-switched communication systems that track and rewrite addresses are not at all unusual. The modern *Multi-Level Packet Switching* (MLPS) technology, which has seen growing use among ISPs and operators of high-performance networks, uses such an approach. Such an approach was also used by ISDN networks and by ATM networks generally, including the backbone network of the telephone system itself.

As for firewalls, these are typically implemented by consumer routers that are connected to ISP networks [2]. Only in exceptional cases would typical ISPs attempt to implement a firewall in

with the aid of other edge services.

[2] In some cases ISPs provide their customers with such devices, butIbelieve this is orthogonal to the question of whether ISPs provide a communication service.

their networks. For one, the complexity of an ISP's network, with numerous interconnections with other networks, would make a firewall difficult to manage. It would also introduce higher latency as routers would inspect each packet to decide whether or not firewall rules should be applied, leading to worse performance generally.

The fact that firewalls will typically be implemented on edge devices connected to the network, or perhaps on top-level routers of ISP customers, makes the wisdom of ISP firewalls questionable. It would appear to be a violation of one of the basic engineering principles of the Internet, the End-to-End Principle, which I describe in Section 4.

3.3.1 Changes in the Form of Information

Relevant to the question of inter-protocol routing and firewalls, the NPRM requests comment on what qualifies as a "change in the form" information. Iclaim, and believe the vast majority of engineers would agree, that the "information" being transmitted over the Internet is in the so-called "payload" of the packets. The "header" information, which includes IP addresses and higher-level information such as "port" numbers, is merely a specification of the end points the users intend to communicate with or among.

Also relevant to this question is the fact that, unknown to and beyond the control of an ISP or its customers, it is possible for other networks to change header information. For example, it is possible that an intermediate network only supports IPv6 routing internally, and simply translates IPv4 addresses at its interconnection points; there are even standards for doing so [2]. Such techniques are merely a part of the operation of the network and should be viewed as incidental to the communication service provided by the network operator.

When ISPs have attempted to modify the content being communicated across their networks it has typically been harmful to consumers. For example, in 2010, the DSL provider Windstream admitted to interfering with its customers' use of search engines, replacing their customers' intended query with results from Windstream's own search engine in an effort to gain additional revenue [5]. In 2011 it was revealed that a company called Paxfire had conspired with several ISPs to engage in similarly abusive practices [4]. Ibelieve these two examples answer the NPRM's request for comment on how consumers were harmed under the weaker regulatory regime that preceded the Title II order.

4 The End-to-End Principle

Key to the Internet's rapid success and displacement of other communication systems is the lack of coordination needed to introduce new applications or to extend the network. This is a consequence of a specific design principle known as the *end-to-end* principle [9]. At a high level the principle states that the network should do only as much as is needed to ensure packets are delivered to their destination address, and only on a best-effort basis. Other, more advanced functionality should, where it is possible to do so, be handled by the edge computers.

As an example, it is often desirable in a packet-switched network to simulate the sort of end-to-end connection available in a circuit-switched network such as the telephone system. It is possible

to design a network that provides such a feature, and standards exist for doing so. Also possible is for the edge computers to simulate a connection by exchanging packets, and likewise standards exist for this approach.

Simulating a connection at the edges is preferred according to the end-to-end principle, and despite years of effort to deviate from this approach [1] it remains dominant on the Internet. This is due to the technical superiority of handling such features at the network's edges. It greatly simplifies the design of the network, which simultaneously reduces the cost of building and operating the network as well as improving the overall performance [12].

4.1 Prioritization

Often it is claimed that violations of the end-to-end principle are sometimes necessary. Such a claim is made in the NPRM, which in seeking comment on the "No Throttling" rule states that differentiated services are beneficial to consumers. The idea behind such claims is that some applications require lower latency, while others require higher throughput, and that ISPs can better serve users by using different routing strategies for different applications.

Unfortunately the neat picture painted by proponents of differentiated services typically fails in practice. It is not generally possible for an ISP to distinguish between classes of applications. In their analysis, Riley and Topolski give as an example the situation of video streaming applications in 2005 [8]. At the time it was common for video streaming to occur over special-purpose protocols. If an ISP had assigned a lower-latency class of service to streaming protocols and a higher-throughput class to the Web, it would have been harmful to Youtube, which uses HTTP as a protocol for streaming video. Had such practices been widespread when Youtube first launched, it might not have grown into one of the most popular streaming video services on the Internet.

5 Innovation at the Edges

The NPRM's characterization of the harm caused by Title II regulation is apparently based on a misunderstanding of where innovation occurs on the Internet. The primary concern of the NPRM is that Title II regulation hampers efforts on the part of ISPs to make innovative offerings to their customers. Yet history tells a much different story of Internet innovations, one in which ISPs have played only a minor role.

From the very beginning, the bulk of Internet innovation occurred at the edges. At first the mapping of human-readable domain names to IP addresses was distributed to all connected computers, like a phone book. Then a new edge service, DNS, was developed to better handle the growing number of edge services requiring human-readable names, a change that required no assistance from network operators. Likewise, the early Gopher system for hypertext documents was replaced by HTTP, without any intervention or assistance on the part of ISPs. HTTP became the "killer application" of the Internet, and, without the assistance of ISPs, HTTP and its hypertext standard HTML have been updated numerous times to support the needs of growing numbers of websites and users. The "Web 2.0" applications of today were developed *at the edges* of the Internet, not by ISPs and without the prompting of ISPs.

Innovation has not been limited to edge *services* i.e. applications with clean divisions between producers and consumers. Peer-to-peer applications allow users to connect directly to each other, without relying on a service provider. Many Internet users rely on peer-to-peer applications, with the BitTorrent application remaining one of the most popular methods of distributing large files. Peer-to-peer applications are often more reliable than client-server applications because there is no single point of failure, which is one of the motivations behind the design of Bitcoin, a peer-to-peer "cryptocurrency" that has attracted enormous interest and investment [6].

When ISPs have attempted to "innovate" the results have typically been poor. The abusive practice on the part of Comcast in 2005 is illustrative [10]. Seeing its network performance degraded by the BitTorrent application, Comcast attempted to remedy the situation by deploying a new system that interfered with BitTorrent. Besides harming its own customers, Comcast's approach was suboptimal and was quickly rendered obsolete by an innovation at the edge: a new version of BitTorrent's protocol that effectively reduces BitTorrent's impact on other applications [7].

6 ISP Services

The central issue of the net neutrality debate is whether ISPs provide a communication service or an information service. The NPRM makes certain assertions about this issue while requesting specific comment on particular details. In the NPRM's analysis, ISPs provide their customers with the capability to access, store, and process information.

Unfortunately the NPRM's analysis is incorrect and presents a confused understanding of the Internet. The truth is that edge services provide an information processing capability, while ISPs facilitate the communication needed to interact with those edge services. It is sometimes the case that ISPs, in addition to Internet connection service, provide edge services of their own, but these are secondary at best, and more typically superfluous.

6.1 DNS

One possible exception, which the NPRM seeks specific comment on, is the DNS service typically provided by ISPs. DNS services map a human-readable "domain name" to a numeric IP address, allowing applications to present a simpler interface to users. The NPRM incorrectly asserts that DNS service is fundamentally part of the package of services ISPs provide, and requests comment on how a lack of such service from ISPs would impact the Internet.

In reality ISPs are not the sole providers of DNS service, and Internet users can easily use a different DNS server. DNS is just another kind of edge service that can be provided by any Internet-connected computer. If ISPs did not commonly provide DNS service the most likely outcome would be consumer devices being sold with pre-configured third-party DNS servers. Moreover, it is possible for ISPs today to configure their customers' equipment to use third-party DNS servers, avoiding the need to operate their own DNS servers. From the perspective of Internet users there would be no difference at all.

6.2 Caching

Another ISP service the NPRM requests specific comment on is caching, whichIpresume refers to HTTP caching in particular. This is a somewhat complex issue, as it may actually refer to different combinations of services provided by ISPs and arrangements made between ISPs and edge service providers. Ibelieve it is appropriate to separately considered different services and arrangements in deciding how to regulate ISPs.

ISP Caching Though more common historically, ISPs may perform caching services on behalf of their users. Such services require an ISP to inspect its users' communication with an edge service and, when possible, intercept that communication with a previously observed response from the edge service. This form of caching is a trade-off for users, potentially speeding up the loading of web pages but also potentially resulting in outdated copies of a page being retrieved.

Importantly, changes in technology meant to improve computer security now hamper this form of caching. Ever-growing numbers of websites encrypt their communication with users, making direct inspection difficult. Transparent caching is also hampered by the increasingly rich and interactive nature of web pages, the growth of social networking features in web applications, and other changes in the edge services market. Transparent caching has also become less important as the capacity of ISP networks and of the Internet generally have grown, to the point where transparent caches can actually worsen performance by increasing latency. This suffices to answer the NPRM's specific question about how the end of caching might affect users; at worst it would not, but most likely in the long-term it would benefit users.

Co-located Edge Provider Caches An increasingly common practice is for edge services to co-locate their own servers with ISP routers to reduce the number of "hops" between users and the service. This should be considered orthogonal to both the classification of ISPs and to net neutrality regulations. In particular, this represents an edge provider connecting its own equipment directly to an ISP network; in other words the edge provider has become a customer of the ISP. ISPs should allow their customers to communicate directly with one another and indeed they do so when they provide their customers with one or more public IP addresses.

7 ISP Abuses prior to Net Neutrality Rules

The NPRM asks for specific comment on what consumer harms were evidence prior to the FCC's efforts to regulate network neutrality, and does admit at least several isolated examples exist. Inote that there was only a brief period during which both the physical infrastructure and routers of ISPs were regulated exclusively under Title I authority. Ialso note that during that same period of time consumer rights groups were publicly demanding net neutrality regulation, and that Congressional hearings were held on possible legislation to that effect. These various factors may have limited abuses somewhat.

That said, there is no question that ISPs violate net neutrality principles repeatedly, and even now, with Title II regulations in place, ISPs continue to violate net neutrality in principle. In the years prior to Title II classification, Verizon sued the FCC over the Open Internet Order and repeatedly testified in court that they wished to experiment with a paid-prioritization business model.

Comcast deliberately interfered with its customers' BitTorrent usage. A small ISP even went as far as replacing the advertisements in web pages its customers loaded with its own paid ads.

Even following the reclassification ISPs have sought ways to violate net neutrality. An oversight in the current framework is the allowance for utilization caps to be selectively applied, so that users are billed differently depending on which applications or edge services they use. In many cases, once the cap is reached, a user's connection with be throttled; it is not yet clear if ISPs that have deployed so-called "zero-rating" would also exempt certain services or applications from such throttling.

The harm to consumers of net neutrality violations is both immediate and long-term. Immediately, net neutrality violations prevent consumers from using the applications and services of their choosing, hampering education, employment, and their ability to interact with government services. The longer-term harm is more troubling, however, as net neutrality violations hamper efforts to develop and deploy new applications.

ISPs that violate net neutrality may also worsen their users' computer security by interfering with the use of encrypted communication services, such as VPNs. The government of China, for example, blocks VPNs, Tor, and other security technologies as part of its censorship program. One troubling possibility is that an ISP might require its customers to pay a premium to use such technologies; such practices are not unheard of in other countries, and American ISPs have charged customers a premium for opting-out of having the ISP record and monetize their web browsing history.

8 Conclusion

The paramount important of network neutrality as a principle of ISP service is beyond dispute. It is one of the basic engineering principle of the Internet, and it is critical to the Internet's ability to support *general* applications. Adherence to net neutrality allowed the vibrant ecosystem and market for Internet applications to develop and to continue developing. Innovation on the Internet is a result of network neutrality rules.

Contrary to Chairman Pai's public statements on this matter, the facts do not support any proposal to do away with network neutrality rules, and by extension the facts do not supoprt any plan to reclassify ISPs under Title I. For good reason the general public does not trust ISPs, unconstrained by market forces due to their de facto monopolies, to voluntarily adhere to network neutrality as a principle. Long before the current regulation ISPs were abusing their de facto monopolies, which existed due to a previous effort at deregulation by the FCC. If the FCC has a plan to resolve the current market failure, let the FCC enact that plan first, and when there is a healthy broadband market deregulation might be considered.

Classifying ISPs as Title II telecommunication services was the right move and represented a job well done by the FCC. I hope the FCC will continue to do its job and continue to appropriately regulate ISPs under Title II. The FCC has a duty to the American public to keep the Internet free and open by maintaining strong network neutrality rules; I hope the FCC will not abdicate its duty.

References

[1] R. Braden, L Zhang, S. Berson, S. Herzog, and C. Jamin. Rfc2205: Resource reservation protocol (rsvp) – version 1 functional specification, 1997. (Cited on page 7.)

[2] A. Conta and S. Deering. Rfc2473: Generic packet tunneling in ipv6 specification, 1998. (Cited on pages 5 and 6.)

[3] IETF Network Working Group. Rfc1122: Requirements for internet hosts – communication layers, 1989. (Cited on page 3.)

[4] Christian Kreibach, Nicholas Weaver, Vern Paxson, Peter Eckersley, and Cindy Cohn. An update on paxfire and search redirection. https://www.eff.org/deeplinks/2011/08/update-paxfire-and-search-redirection, 2011. (Cited on page 6.)

[5] Matthew Lasar. Windstream in windstorm over isps search redirects. https://arstechnica.com/tech-policy/2010/04/windstream-in-windstorm-over-dns-redirects/, 2010. (Cited on page 6.)

[6] Satoshi Nakamoto. Bitcoin: A peer-to-peer electronic cash system. http://bitcoin.org/bitcoin.pdf, 2008. (Cited on page 8.)

[7] Arvid Norberg. utorrent transport protocol. bittorrent.org/beps/bep_0029.html, 2009. (Cited on page 8.)

[8] Chris Riley and Robb Topolski. The hidden harms of application bias. *New America Foundation and Free Press*, 2009. (Cited on page 7.)

[9] J. H. Saltzer, D. P. Reed, and D. D. Clark. End-to-end arguments in system design. *ACM Trans. Comput. Syst.*, 2(4):277–288, November 1984. (Cited on page 6.)

[10] Ryan Singel. Comcast sued over bittorrent blocking - updated. https://www.wired.com/2007/11/comcast-sued-ov/, 2007. (Cited on page 8.)

[11] P. Sriruresh and M. Holdrege. Rfc 2663: Ip network address translator (nat) terminology and considerations, 1999. (Cited on page 5.)

[12] Internet2 Vice President Testimony of Gary R. Bachula. Network neutrality: Preserving openness and innovation on the internet. (Cited on pages 2 and 7.)